arXiv:1610.03564v1 [cs.GT] 11 Oct 2016

Finding Bidder-Optimal Core Points Quickly

Jason Hartline*[1], Nicole Immorlica†[2], Brendan Lucier‡[2], and Rad Niazadeh§[3]

[1]Department of Computer Science, Northwestern University.
[2]Microsoft Research, New England.
[3]Department of Computer Science, Cornell University.

October 13, 2016

Abstract

In complex combinatorial markets with complementary valuations, truthful auctions can yield low revenue. Core-selecting auctions attempt to boost revenue by setting prices so that no group of agents, including the auctioneer, can jointly improve their utilities by switching to a different allocation and payments. Among outcomes in the core, bidder-optimal core points have been the most widely studied due to their incentive properties, such as being implementable at equilibrium. Prior work in economics has studied heuristics and algorithms for computing approximate bidder-optimal core points given oracle access to the welfare optimization problem, but these solutions either lack performance guarantees or are based on prohibitively convex programs. Our main result is a combinatorial algorithm that finds an approximate bidder-optimal core point in a quasi-linear number of calls to the welfare maximization oracle. Our results have implication for the sale of space in online advertising.

*hartline@eecs.northwestern.edu
†nicimm@microsoft.com
‡brlucier@microsoft.com
§rad@cs.cornell.edu

1 Introduction

Auctions with combinatorial preferences are prevalent in practice. Prominent examples include the sale of wireless spectrum and the sale of online advertising space. In the former, companies have combinatorial and potentially complementary preferences over wireless bands in various geographical regions. In the later, companies with ads of various shapes and configurations have combinatorial and potentially complementary preferences over ad space on search engines and online content pages like the New York Times. In both settings, sellers use auctions to determine the allocation and payments.

There are many potential auction formats for combinatorial markets, among which there are two of particular note: Vickrey-Clark-Groves (VCG) auctions and core-selecting auctions. These auctions are direct revelation mechanisms, that is, they ask bidders to report their preferences over bundles of items and then compute an allocation and payments as a function of these reports. Both auctions select the optimal welfare allocation, i.e., the one with the maximum total value for the bidders. Optimizing welfare is a computationally hard problem in many cases of interest, and this issue has drawn considerable attention from the algorithmic game theory community. However, even in settings where this difficulty can be resolved satisfactorily (e.g., via heuristics that have exponential runtime in the worst case but tend to solve practical instances quickly), this still leaves the problem of computing the payments.

In VCG autions, the payment of a bidder is, intuitively, the harm that he imposes on the other bidders by consuming the resources allocated to him. While rarely used in practice [3], this auction is a useful point of comparison as it has several important theoretical properties. First, among welfare-maximizing auction formats, it is the unique payment rule (up to additive offsets) that incentivizes truthtelling as a dominant strategy on the part of the bidders. Second, given access to an oracle which computes the optimal welfare allocation, it is computationally efficient: it takes just $n + 1$ oracle calls to compute the outcome (allocation and payments).

Unfortunately, in the presence of complementarities, the revenue generated by VCG can be quite low compared to the bidders' values, and the resulting outcome can seem unfair. Consider, for example, a setting with two items, A and B, and three bidders, 1, 2, and 3. Suppose bidder 1 only wants item A and has a value of \$100 for it. Similarly, bidder 2 only wants item B and also has a value of \$100 for it. Bidder 3 has complementary preferences. He only wants both A and B and has a value of \$199 for this bundle. In this setting, the VCG auction gives item A to bidder 1 and item B to bidder 2 and charges each of them a price of \$1 for a total revenue of \$2. This revenue is both low compared to the values, and also seemingly unfair from the point of bidder 3 who would be willing to pay quite a bit more than the winners.

Core-selecting auctions attempt to address both the revenue and fairness issues of VCG while retaining welfare optimality and minimally sacrificing the incentives.[1] They again choose an optimal welfare allocation, but then set payments such that no group of bidders, including the auctioneer, can simultaneously improve outcomes by switching to a different allocation and payments. In the above example, if the bidders bid their true values, then any set of payments such that bidders 1 and 2 jointly pay at least \$199 while each paying at most \$100 could be the outcome of a core-selecting auction.

Core-selecting auctions are not truthful in the sense of dominant strategies, but they have natural equilibria that generate welfare-optimal outcomes [12, 10]. Most of the work on core-selecting auctions has focused on bidder-optimal payments, which have the property that it is not possible to reduce any bidders payment and remain in the core. In the above example, having each of bidders 1 and 2 pay \$100 is a core outcome, but is not bidder-optimal; having bidder 1 pay \$100 and bidder 2 pay \$99 *is* bidder-optimal. Bidder-optimal core payments have additional desirable incentive properties; see [9]

[1] Another common way to increase revenue is to impose reservation prices. However, reserve prices reduce welfare by design and are necessarily prior-dependent.

for a discussion. As a result, they have been used in practice to sell wireless spectra [8] and have been proposed for use in selling online advertising [10, 16].

However, to be truly practical for large frequent auctions like online advertising, an auction must be highly computationally efficient. As core-selecting auctions output optimal-welfare allocations, the computation associated with finding such allocations is unavoidable. The main issue we tackle in this paper is the additional difficulty of computing core payments, given a satisfactory solution to the welfare optimization problem. To make this separation clear, we make the standard assumption of oracle access to a welfare optimization algorithm. Prior work on core payments assumes access to such an oracle and computes core outcomes using either heuristics [10, 12, 14, 6] or computationally intensive convex programming methods [12].

Our Results Our main result is a quasi-linear time algorithm for computing bidder-optimal core payments, given access to an oracle that finds a welfare-optimal allocation for a given profile of buyer valuations.[2]

Theorem (informal): There is an algorithm that computes an ϵ-approximate bidder optimal core point, using $O(n \log(n/\epsilon))$ calls to a welfare-optimization oracle.

The set of core payoffs is a polytope, determined by the (exponentially many) constraints that no subset of bidders can simultaneously improve their utilities. Intuitively, the algorithm proceeds by exploring this polytope. Starting from an arbitrary point in the core, the algorithm attempts to increase bidder utilities, which corresponds geometrically to following a positively-oriented ray until hitting a facet of the polytope. It then determines whether there exists a subset of the bidders whose utilities can be increased while remaining in the core, and if so continues to follow an appropriate ray. This process repeats until a bidder-optimal core point is reached.

This geometric intuition corresponds to a water-filling algorithm, in which the utilities of agents are simultaneously increased and frozen as constraints become tight. Implementing this approach requires an efficient test for inclusion in the core polytope, as well as a method for determining which bidders are involved in a tight core constraint. It turns out that both questions can be expressed as queries to the welfare-optimization oracle, with appropriately truncated valuation profiles. This is the most technical contribution of the paper, and involves an analysis of the geometry of the core constraints. This is naturally related to the well-known equivalence of optimization and separation, and indeed one can interpret our results as exploiting the geometry of the core to implement and employ a specially-tailored separation oracle. Given these tests, finding a tight core constraint can be done using binary search (which is what generates the logarithmic factors in our runtime). The algorithm terminates after at most n iterations, since each will freeze the payment of at least one bidder.

Compared to the VCG auction, our algorithm finds payments nearly as quickly (with quasi-linear as opposed to linear number of oracle calls), and produces higher revenue and a more fair outcome on the same profile of bids. But as bidders are strategic, any comparison must consider the outcomes at equilibria of the auctions. The computational properties are worst-case and hence do not degrade with strategic play. But what about the welfare, revenue, and fairness properties? As we discuss in Section 5, our auction has a natural ϵ-Nash equilibrium in which players play truncation strategies (i.e., shave all values by an additive constant). This equilibrium is welfare-optimal and has higher revenue than VCG's dominant strategy equilibrium. Furthermore, by varying the direction of binary search in the core polytope, our algorithm can be parameterized to favor different core outcomes. For example, we can maximize the minimum utility enjoyed by any winner (subject to being in the core), or we can attempt to equalize, across all winners, the ratio between their utility and the utility they would obtain in the VCG outcome.

[2] Technically, we only require that the oracle returns the values of the bidders in a welfare-optimal outcome, and not the allocation itself.

Related Work The VCG mechanism is a prevalent auction for combinatorial settings (for primary sources see Vickrey [24], Clarke [7], Groves [19]), but it is rarely used in practice due to practical issues, including its potentially low revenue; e.g. see Ausubel and Milgrom [2, 3] for a survey. The *ascending proxy auction* was one auction proposed by Ausubel and Milgrom [2] as a practical alternative that resolves some of these issues. This ascending auction terminates at a desirable "bidder Pareto-optimal core outcome." This line of research was further developed by Day and Raghavan [12], Day and Milgrom [10, 11], Day and Cramton [9] by considering the general notion of core-selecting package auctions, and identifying incentive properties of Pareto optimal frontier of core polytope, as well as revenue advantages of core auctions.

Pareto bidder optimal core outcomes are implementable at equilibrium, but are not truthful. Payment rules that pick particular points in the core have been proposed in the literature to mitigate this issue. For example, the difference between a final payment and the VCG payment represents a measure of "residual incentive to misreport," as observed by Parkes et al. [21]. Day and Raghavan [12] therefore proposed the minimum-revenue point and Day and Cramton [9] proposed the closest-to-VCG point, which are the Pareto optimal core points that minimize ℓ_1-norm and ℓ_2-norm of this difference respectively. An alternative practical core payment rule proposed by Erdil and Klemperer [14] considers robustness with respect to the submitted bids.

It is known that computing core points requires black-box access to a winner determination oracle, which is effectively a separation oracle for the core polytope [9]. One can therefore optimize convex objectives over core polytope using the Ellipsoid algorithm [18], but this is rather slow in practice. Various heuristic algorithms have been proposed for linear objectives of revenue (i.e., the *Core Constraint Generation algorithm*) [12, 6], or the quadratic objective of distance to VCG [9, 14].

VCG and core auctions have been compared along multiple dimensions. The Bayesian setting was studied by Ausubel and Baranov [1], and more recently by Sano [23] and Goeree and Lien [17], providing theoretical and experimental evidence that the revenues and efficiency from core-selecting auctions improve as correlations among bidders' values increase, while the revenues from the VCG auction worsen. Another issue with VCG is the lack of revenue monotonicity: adding bidders or increasing bids can potentially decrease the seller's revenue [15, 20, 22, 5]. Beck and Ott [5] and Lamy [20] proposed revenue-monotone core selection.

2 Notations and Basics

We consider a combinatorial auction setting with n bidders and m items. Let N be the set of bidders. The auction asks each bidder to declare a valuation function, which assigns a value to each subset of the items. We will write b_i for the valuation function submitted by agent i. We will assume that valuation functions are normalized so that $b_i(\emptyset) = 0$, and all values are scaled to lie in $[0, 1]$. An allocation is an assignment of item bundles to agents, $\{x_i\}_{i\in N}$, such that $x_i \cap x_j = \emptyset$ for all $i \neq j$.

Given the bid functions $\{b_i\}_{i\in N}$, the auction will return an *outcome*, consisting of an allocation $\{x_i\}_{i\in N}$ and a payment $p_i \geq 0$ for each bidder. We will tend to write π_i for the resulting utility[3] for bidder i, so that $\pi_i = b_i(x_i) - p_i$. We also write π_0 for the revenue of the seller.

Throughout this paper, we assume that the auction's allocation rule does not allocate sets of zero value. We say that a feasible allocation $\{x_i^*\}_{i\in N}$ is welfare-maximizing if and only if $\{x_i^*\}_{i\in N} \in \operatorname{argmax}_{x: x_i\cap x_j=\emptyset, i\neq j} \sum_{i\in N} b_i(x_i)$. For any feasible allocation $\{x_i\}_{i\in N}$, let $W(\{x_i\}_{i\in N}) \triangleq \{i \in N : x_i \neq \emptyset\}$ be the corresponding winner set. Let $w(S, \{b_i(.)\}_{i\in N})$ be the maximum welfare of coalition S with respect to the bids $\{b_i\}$. That is, $w(S, \{b_i(.)\}_{i\in N})$ is the maximum welfare of any allocation that assigns items

[3]Since our focus is the computational problem of finding core payments, and not buyer incentives, we will not differentiate between true and declared valuations in our notation. We abuse notation for convenience and define the terms "utility" and "welfare" with respect to the declared bid functions. We discuss incentive issues in Section 5.

only to agents in S. Unless noted otherwise, we will write $w(S) \triangleq w(S, \{b_i(.)\}_{i\in N})$ for notational convenience.

Definition 2.1 (**Winner Set**). *A subset $W \subseteq N$ of bidders is a winner set with respect to bids $\{b_i(.)\}_{i\in N}$ if and only if there exists a feasible allocation $\{x_i^*\}_{i\in N}$ such that*

$$\{x_i^*\}_{i\in N} \in \underset{x: x_i \cap x_j = \emptyset, i\neq j}{\operatorname{argmax}} \sum_{i\in N} b_i(x_i) \tag{1}$$

and $W = \{i \in N : x_i^* \neq \emptyset\}$. Let $\mathcal{W}(\{b_i(.)\}_{i\in N})$ be the set of winner sets with respect to $\{b_i(.)\}_{i\in N}$.

Definition 2.2 (**Core**). *A vector of non-negative utilities $\{\pi_i\}_{i\in N\cup\{0\}}$ is said to be in the core with respect to the submitted bids $\{b_i(.)\}_{i\in N}$ if no blocking coalition exists, i.e.*

$$\forall S \subseteq N : \pi_0 \geq w(S) - \sum_{i\in S} \pi_i \tag{2}$$

Similarly, an outcome $\{(x_i, p_i)\}_{i\in N}$ is said to be in the core if its corresponding vector of utilities (i.e. $\forall i \in N : \pi_i = b_i(x_i) - p_i$ and $\pi_0 = \sum_{i\in N} p_i$) is in the core.

It is easy to see that any core outcome $\{(x_i, p_i)\}_{i\in N}$ is welfare-maximizing, as for coalition N we have $\sum_{i\in N} b_i(x_i) = \pi_0 + \sum_{i\in N} \pi_i \geq w(N)$. Therefore, one can rewrite Equation 2 as

$$\forall S \subseteq N : w(N) - \sum_{i\in N} \pi_i \geq w(S) - \sum_{i\in S} \pi_i. \tag{3}$$

When clear from context, we will sometimes abuse notation slightly and use a set of bidders S to denote the associated core constraint (3).

The following lemma shows that given a core point in utility space, it is possible to reconstruct core payments. We will therefore focus on the problem of computing core points in utility space.

Lemma 2.3. *Given a core point $\{\pi_i\}_{i\in N}$ and any maximum welfare allocation $\{x_i^*\}_{i\in N}$, let $p_i = b_i(x_i^*) - \pi_i$ for all $i \in N$. Then $\{(x_i^*, p_i^*)\}_{i\in N}$ will be an outcome in the core.*

Proof. By Definition 2.2, as long as payments are non-negative the outcome $\{(x_i^*, p_i^*)\}_{i\in N}$ will be in the core. So, we only need to show for every $i \in N$, $\pi_i \leq b_i(x_i^*)$. By looking at the core constraint for coalition $N \setminus \{i\}$, we have:

$$w(N) - \sum_{j\in N} \pi_j \geq w(N \setminus \{i\}) - \sum_{j\in N\setminus\{i\}} \pi_j \Rightarrow$$
$$\pi_i \leq w(N) - w(N \setminus \{i\}) = \sum_{j\in N} b_j(x_j^*) - w(N \setminus \{i\}) \leq \sum_{j\in N} b_j(x_j^*) - \sum_{j\in N\setminus\{i\}} b_j(x_j^*) = b_i(x_i^*).$$

where the last inequality holds because $\{x_j^*\}_{j\in N\setminus\{i\}}$ is a feasible allocation for bidders in $N \setminus \{i\}$. □

Our goal is to compute bidder optimal core points, i.e., points in the core that are not dominated by any other point in the core with respect to bidder utilities. We will relax this condition by defining ϵ-bidder optimal core points, i.e. core points that are ϵ-close to being bidder Pareto optimal.

Definition 2.4. *A vector of utilities $\{\pi_i\}_{i\in N}$ is said to be ϵ-bidder optimal for some $\epsilon > 0$, if*

- *$\{\pi_i\}_{i\in N}$ is in the core with respect to the submitted bids $\{b_i(.)\}_{i\in N}$, and*
- *there exists no other core point $\{\pi'_i\}_{i\in N}$ such that for every bidder $i \in N$, $\pi'_i \geq \pi_i$, and for at least one bidder j, $\pi'_j > \pi_j + \epsilon$.*

Similarly, an outcome $\{(x_i, p_i)\}_{i\in N}$ is said to be ϵ-bidder optimal if its corresponding vector of utilities is ϵ-bidder optimal.

Observe that the set of core points forms a polytope in utility space, described by the (exponentially many) constraints of the form (3). We are interested in identifying combinatorial structures of the core polytope, in order to design a fast algorithm for finding a bidder optimal point. For this purpose, as we will describe in more detail in later sections, it will be useful to find the core constraint in Definition 2.2 whose right-hand side takes the maximum value. Therefore, we define the notion of *maximum binding core constraints.*

Definition 2.5. *Given a vector of utilities $\{\pi_i\}_{i\in N}$, a subset $S \subseteq N$ is a Maximum Binding Core Constraint (MBCC) if and only if $S \in \operatorname{argmax}_{S'\subseteq N} w(S') - \sum_{i\in S'} \pi_i$.*

Definition 2.6. *Given a vector of utilities $\{\pi_i\}_{i\in N}$, a subset $S \subseteq N$ is an ϵ-tight core constraint (or a tight core constraint when $\epsilon = 0$) if and only if*

$$w(S) - \sum_{i\in S} \pi_i \leq w(N) - \sum_{i\in N} \pi_i \leq w(S) - \sum_{i\in S} \pi_i + \epsilon. \tag{4}$$

3 The Water-Filling Algorithm

In brief, our algorithm is a water-filling method that starts from an arbitrary point in the core (e.g., pay-your-bid, where all agents receive utility 0) and then at each iteration does the following:

- **Finding a Feasible Direction**: it first finds a subset of bidders such that if we increase all of their utilities uniformly by a small amount, we still remain in the core.
- **Uniform Utility Increase**: It then increases the utilities for those bidders uniformly, until it hits a facet of the core (approximately).
- **Checking Termination Condition/Repeat:** Finally, it checks whether there exists any remaining subset of bidders who can increase their utilities and still remain in the core. If so, the algorithm iterates. Otherwise it terminates.

See Figure 1 for an overview of the algorithm, and its geometric interpretation.

3.1 Winner Determination Oracle Model

In this paper, we focus on the *winner determination oracle model* as our model of computation. More exactly, suppose we have access to a winner determination oracle under sincere strategies, where a sincere strategy is a truncation of $\{b_i(.)\}$ by $\{\pi_i\}$, i.e. $\{\max(b_i(.)-\pi_i, 0)\}$. The oracle accepts submitted bids and the required truncations as input, and then simulates truncated bids to find a maximum welfare allocation. To make this more concrete, we define the oracle `WIN-ORAC`.

Definition 3.1. *`WIN-ORAC` is defined to be a black-box oracle with the following input-output relation:*

- *`Input:` submitted bids $\{b_i(.)\}_{i\in N}$ and required truncations $\{\pi_i\}_{i\in N}$.*
- *`Output`: maximum welfare with truncated bids, i.e. $w(N, \{\max(b_i(.) - \pi_i, 0)\}_{i\in N})$, and a winning set under truncated bids, i.e. $S \in \mathcal{W}(\{\max(b_i(.) - \pi_i, 0)\}_{i\in N})$ (or `NULL` if no such set exists).*

3.2 The Algorithm

Here is the full description of our algorithm (Algorithm 1). The algorithm uses the subroutine `Core-Search` at each iteration, which is basically binary-search to find the next feasible subset of dimensions quickly, and it only requires $\tilde{O}(1)$ number of calls to the oracle `WIN-ORAC`.

The iterations of the algorithm are indexed by t. The algorithm maintains a current core point $\pi^{(t)}$ (initially the origin) and a collection of active bidders, S_t, whose utilities can be increased (initially all bidders). On each iteration, the algorithm updates the set of active bidders S_t by finding and removing a set $S_t \setminus T_t$ of bidders whose utilities cannot be increased without violating a core constraint (Finding the next feasible set might take several iterations, without changing the point $\pi^{(t)}$, for reasons that will become clear in the proof). It then applies binary search along the ray of points consisting of uniform increases to the utilities of all agents in S_{t+1}, described in subroutine `Core-Search`. This finds a pair of points $\{\underline{\pi}_i\}$ and $\{\bar{\pi}_i\}$, where $\underline{\pi}$ is in the core, $\bar{\pi}$ is outside the core, and the points are within ϵ of each other in ℓ_1 distance, i.e.

$$w(N) - \sum_{i \in N} \underline{\pi}_i \triangleq \underline{\pi}_0 \geq w\left(N, \{\max(b_i(.) - \underline{\pi}_i, 0)\}\right) \qquad \triangleright \text{ feasible core point.}$$

$$w(N) - \sum_{i \in N} \bar{\pi}_i \triangleq \bar{\pi}_0 < w\left(N, \{\max(b_i(.) - \bar{\pi}_i, 0)\}\right) \qquad \triangleright \text{ infeasible point out of the core.}$$

$$\sum_{i \in N} \bar{\pi}_i - \sum_{i \in N} \underline{\pi}_i \leq \epsilon$$

The algorithm uses $\underline{\pi}$ as its updated core point, and uses $\bar{\pi}$ to update the set of active bidders in the subsequent iteration, by finding the set of winners for sincere bids $\{\max(b_i(.) - \bar{\pi}_i, 0)\}$, i.e. T_{t+1}. Once all bidders are frozen, the algorithm returns the current core point. An illustration of each iteration and how sets S_t and T_t are set at each iteration can be seen in Figure 1, which describes simulating our algorithm on the example in [12].

Algorithm 1 Water-filling for finding an ϵ-bidder optimal core point

parameters: $\epsilon > 0$.
input: submitted bids $\{b_i(.)\}$, set of bidders N.
initialize: $\pi_i^{(0)} = 0$ and $\bar{\pi}_i^{(0)} = 0$ for all $i \in N$. ▷ utility vector of the pay-your-bid auction.
$t \leftarrow 0$, $S_0 \leftarrow N$.
while $S_t \neq \emptyset$ **do**
    $T_t \leftarrow$ any set in $\mathcal{W}(\{\max(b_i(.) - \bar{\pi}_i^{(t)}, 0)\})$ ▷ needs a query call to `WIN-ORAC`.
    $S_{t+1} \leftarrow S_t \cap T_t$
    if $S_{t+1} \neq \emptyset$ **then**
        Run subroutine `Core-Search`$(\{b_i(.)\}, \pi^{(t)}, S_{t+1})$ to return $\{\bar{\pi}_i\}_{i \in N}$ and $\{\underline{\pi}_i\}_{i \in N}$.
        $\pi_i^{(t+1)} \leftarrow \underline{\pi}_i$, $i \in N$.
        $\bar{\pi}_i^{(t+1)} \leftarrow \bar{\pi}_i$, $i \in N$
        $t \leftarrow t + 1$.
    end if
end while
return $\pi^{(t)}$

4 Proof of Correctness and Running Time

The main idea behind the proof of correctness of the algorithm is the following simple observation. Suppose π is a point in the core and S is a tight core constraint with respect to π. Note that there

Procedure `Core-Search` for water-filling algorithm

parameters: $\epsilon > 0$.
input: submitted bids $\{b_i(.)\}_{i\in N}$, core point π and subset of bidders $S \subseteq N$.
initialize: $\Delta_l = 0$ and $\Delta_h = 1$.
while $\Delta_h - \Delta_l > \frac{\epsilon}{|S|}$ **do** ▷ do binary-search to find Δ_l and Δ_h.
    $\Delta \leftarrow \frac{\Delta_l + \Delta_h}{2}$.
    Let $\tilde{\pi}_i = \pi_i + \Delta$ for $i \in S$, $\tilde{\pi}_i = \pi_i$ for $i \in N \setminus S$, and $\tilde{\pi}_0 = w(N) - \sum_{i\in N} \tilde{\pi}_i$.
    if $\tilde{\pi}_0 \geq w\left(N, \{\max(b_i(.) - \tilde{\pi}_i, 0)\}\right)$ **then** ▷ requires one query call to oracle `WIN-ORAC`.
        $\Delta_l \leftarrow \Delta$.
    else $\Delta_h \leftarrow \Delta$.
    end if
end while
return $\{\bar{\pi}_i\}_{i\in N}, \{\underline{\pi}_i\}_{i\in N}$, where $\begin{cases} \bar{\pi}_i = \pi_i + \Delta_h \text{ and } \underline{\pi}_i = \pi_i + \Delta_l, & \text{if } i \in S. \\ \bar{\pi}_i = \pi_i \text{ and } \underline{\pi}_i = \pi_i, & \text{otherwise.} \end{cases}$

always exists at least one tight core constraint, as the constraint for coalition N is always tight. Now, one is allowed to increase π_i by an small amount and still have a core point as long as bidder i is participating in *every* tight core constraint $S \subseteq N$. This is true because the change in the left-hand side and right-hand of Equation 3 will be the same for all currently tight constraints, which are the only candidates to get violated after the small change. Inspired by this observation, the water-filling algorithm starts from a point in the core and fills the utilities of nested subsets of bidders uniformly at each iteration, until it ensures that there exists no bidder in the intersection of all tight (and almost tight) core-constraints.

Given this observation, here is the intuition behind the correctness of the algorithm. As mentioned earlier, Algorithm 1 keeps track of $\{\pi_i^{(t)}\}_{i\in N}$ (inside the core), and $\{\bar{\pi}_i^{(t)}\}_{i\in N}$ (outside of the core), while these two points are always ϵ-close in ℓ_1-norm distance. As a result, they help the algorithm to find the next subset of bidders $S_{t+1} \subset S_t$, with the property that the algorithm can potentially increase their corresponding utilities uniformly. Furthermore, for a fixed run of the algorithm consider sequence $\emptyset = \mathcal{G}_0 \subset \mathcal{G}_1 \subset \mathcal{G}_2 \subset, \ldots$, where $\mathcal{G}_t \triangleq \{T_1, \ldots, T_{t-1}\}$. By closeness of the two points, it can be shown that $\mathcal{G}_t$ is always a collection of ϵ-tight core constraints with respect to $\pi^{(t)}$. This sequence acts as a *certificate of correctness* for the algorithm: at each iteration t the algorithm ensures that bidders in $S_{t+1} = \bigcap_{t'=0}^{t} T_t$ are indeed the subset of bidders appearing in *every* constraint of $\mathcal{G}_{t+1}$, and hence those which could have increased their utility. So, at termination it ensures that S_{t+1} is empty, and therefore there exists no bidder that appears in all ϵ-tight core constraints.

More concretely, we prove the correctness of the algorithm in two steps. We first show upon termination the algorithm outputs a point in the core that is ϵ-bidder optimal, and then we show that the algorithm terminates in at most $|N| = n$ iterations. To do so, we start by proving the following lemma, which is crucial in understanding how our algorithm works.

Lemma 4.1. *Fix a vector of utilities $\{\pi_i\}_{i\in N}$. Suppose $S \subseteq N$. The following are true:*

- *S is a maximum binding constraint if it is also the set of winning bidders for some maximum welfare allocation under bids $\{\max(b_i(.) - \pi_i, 0)\}_{i\in N}$.*
- *Suppose $\forall i \in S : \pi_i > 0$. Then S is a maximum binding constraint only if it is the set of winning bidders for some maximum welfare allocation under bids $\{\max(b_i(.) - \pi_i, 0)\}_{i\in N}$.*

Proof. Suppose S' be the set of winners for a maximum welfare allocation x' under truncated bids $\{\max(b_i(.) - \pi_i, 0)\}_{i\in N}$. Note for all $i \in S'$, $b_i(x'_i) > \pi_i$ (otherwise, allocation x' gives items for free

to some bidder j, because $x'_j \neq \emptyset$ & $\max(b_j(x'_j) - \pi_j, 0) = 0$ as $b_j(x'_j) \leq \pi_j$). Let S be a maximum binding constraint and let x^S be the maximum welfare allocation restricted to bidders in S. We have

$$\begin{aligned} w(S) - \sum_{i\in S} \pi_i &= \sum_{i\in S} (b_i(x_i^S) - \pi_i) \leq \sum_{i\in N} \max(b_i(x_i^S) - \pi_i, 0) \\ &\leq \sum_{i\in N} \max(b_i(x'_i) - \pi_i, 0) = \sum_{i\in S'} (b_i(x'_i) - \pi_i) \leq w(S') - \sum_{i\in S'} \pi_i \end{aligned} \tag{5}$$

and therefore S' is also a maximum binding constraint.

Next, suppose S is a maximum binding constraint. Let x^S be the allocation that maximizes the welfare restricted to bidders in S. Note that in such an allocation, all the bidders in S will be winners, i.e. $x_i^S \neq \emptyset$ for $i \in S$ (because otherwise there exists $S' \subset S$ such that $w(S) = w(S')$, and therefore $w(S) - \sum_{i\in S} \pi_i < w(S') - \sum_{i\in S'} \pi_i$ which is a contradiction). Now let x' be a maximum welfare allocation for truncated bids, and let S' be the set of winners under such an allocation. We have:

$$\begin{aligned} \sum_{i\in N} \max(b_i(x'_i) - \pi_i, 0) &= \sum_{i\in S'} (b_i(x'_i) - \pi_i) \leq w(S') - \sum_{i\in S'} \pi_i \leq w(S) - \sum_{i\in S} \pi_i \\ &= \sum_{i\in S} (b_i(x_i^S) - \pi_i) \leq \sum_{i\in S} \max(b_i(x_i^S) - \pi_i, 0) = \sum_{i\in N} \max(b_i(x_i^S) - \pi_i, 0). \end{aligned} \tag{6}$$

Therefore x^S is a maximum welfare allocation under $\{\max(b_i(.) - \pi_i, 0)\}_{i\in N}$, and S is its winner set. □

Now, using Lemma 4.1, we can prove the following invariants of our algorithm: $\pi^{(t)}$ is a core point for each t, and $\{\mathcal{G}_t\}$ is a collection of ϵ-tight core constraints with respect to $\pi^{(t)}$.

Lemma 4.2. *Given submitted bids $\{b_i(.)\}_{i\in N}$ and $\epsilon > 0$, there exists a finite sequence of collections $\emptyset = \mathcal{G}_0 \subset \mathcal{G}_1 \subset \ldots \subset \mathcal{G}_{T+1}$, such that the following invariants hold at each iteration t of Algorithm 1:*

(1) $\{\pi_i^{(t)}\}_{i\in N}$ is always in the core, and $\{\bar{\pi}_i^{(t)}\}_{i\in N}$ is outside of the core for $t \geq 1$.

(2) S_t is the subset of bidders that are simultaneously participating in all core constraints included in the collection $\mathcal{G}_t \triangleq \{T_0, \ldots, T_{t-1}\}$. Moreover, for $t \geq 1$, $S_t \setminus T_t \neq \emptyset$ and $\mathcal{G}_t \neq \mathcal{G}_{t+1}$.

(3) $\mathcal{G}_{t+1}$ is a collection of ϵ-tight core constraints with respect to $\pi^{(t)}$ (as in Definition 2.6).

Proof.
Part (1): To prove this part, we use induction. For the base case $t = 0$, all-zero vector $\pi^{(0)}$ is always in the core as $w(N) \geq w(S)$ for all $S \subseteq N$. Now suppose $\pi^{(t-1)}$ is in the core. If $\pi^{(t)} = \pi^{(t-1)}$ we are done. So let $\pi^{(t)} \neq \pi^{(t-1)}$. As a result, $\pi^{(t)} = \underline{\pi}$ in the binary search phase at iteration $t-1$. Therefore, due to termination condition of `Core-Search` subroutine, for all $S \subseteq N$ we have:

$$w(N) - \sum_{i\in N} \pi_i^{(t)} = \underline{\pi}_0 \geq w\left(N, \{\max(b_i(.) - \pi_i^{(t)}, 0)\}\right) \geq w(S) - \sum_{i\in S} \pi_i^{(t)} \tag{7}$$

where the last inequity holds because of Lemma 4.1. So, $\pi^{(t)}$ is a core point. Also, for $t \geq 1$ the point $\bar{\pi}^{(t)}$ is outside of the core, because once binary search stops at iteration $t-1$ we have:

$$w(N) - \sum_{i\in N} \bar{\pi}_i^{(t)} = \bar{\pi}_0 < w\left(N, \{\max(b_i(.) - \bar{\pi}_i^{(t)}, 0)\}\right) = w(S) - \sum_{i\in S} \bar{\pi}_i^{(t)} \tag{8}$$

for some $S \subseteq N$.
Part (2): First of all, note that S_t is the subset of bidders that are simultaneously participating in all of $\{T_{t'}\}_{t=1}^{t-1}$, simply because $S_t = \bigcap_{t'=0}^{t-1} T_{t'}$ due to the update rule of S_t. Moreover, at iteration $t-1$ (for

$t \geq 1$) the algorithm starts from a feasible core point $\pi^{(t-1)}$ and uniformly increases the utilities only for bidders in S_t, until it reaches to a point $\bar{\pi}^{(t)}$ that is outside of the core. Note that no constraint $S \supseteq S_t$ will get violated during this process, simply because the changes in the left hand side and right hand side of these constraints (refer to Definition 2.2 and Equation 3) are equal, following the fact that $S_t = \bigcap_{t'=0}^{t-1} T_{t'}$. In particular, no constraint in $\mathcal{G}_t$ will get violated by $\bar{\pi}^{(t)}$. At iteration t, T_t is set to one of the most binding core constraint with respect to $\bar{\pi}^{(t)}$ (due to Lemma 4.1), and therefore it should be a violated core constraint, because $\bar{\pi}^{(t)}$ is outside of the core. Combining the above arguments, $S_t \setminus T_t \neq \emptyset$. So, $T_t \notin \mathcal{G}_t$ and therefore $\mathcal{G}_t \neq \mathcal{G}_{t+1}$.

Part (3): To prove this part, we again use induction. For the base case $t = 1$, $\mathcal{G}_1 = \{T_0\}$ is the set of winning bidders of maximum welfare allocation under original bids $\{b_i(.)\}_{i \in N}$. Therefore, due to Lemma 4.1, it is also a maximum binding constraint for all-zero vector $\pi^{(0)}$. Moreover, coalition N is always a tight core constraint (as in Definition 2.6), and therefore any maximum binding core constraint is also a tight core constraint. Now suppose $\mathcal{G}_t = \{T_0, \ldots, T_{t-1}\}$ is a collection of ϵ-tight core constraints with respect to $\pi^{(t-1)}$. Note that if a constraint is ϵ-tight core constraint at some iteration, it will always remain ϵ-tight, as utilities never decrease and $\pi^{(t)}$ is also in the core. At iteration t, $\pi^{(t)} = \underline{\pi}$ and $\bar{\pi}^{(t)} = \bar{\pi}$, where these parameters are set in binary search phase at the previous iteration $t-1$. Moreover, T_t will be a violated constraint under $\bar{\pi}$ (as in the proof of Part (2)). Therefore,

$$w(N) - \sum_{i \in N} \pi_i^{(t)} \overset{(1)}{\leq} w(N) - \sum_{i \in N} \bar{\pi}_i^{(t)} + \epsilon \overset{(2)}{<} w(T_t) - \sum_{i \in T_t} \bar{\pi}_i + \epsilon \overset{(3)}{\leq} w(T_t) - \sum_{i \in T_t} \pi_i^{(t)} + \epsilon \tag{9}$$

where inequality (1) holds because $\sum_{i \in N} \bar{\pi}_i - \sum_{i \in N} \underline{\pi}_i \leq \epsilon$, inequality (2) holds because T_t is a violated core constraint for $\bar{\pi}^{(t)} = \bar{\pi}$, and inequity (3) simply holds because for all $i \in N$, $\pi_i^{(t)} = \underline{\pi}_i \leq \bar{\pi}_i$. So, T_t is also an ϵ-tight core constraint for $\pi^{(t)}$, and therefore $\mathcal{G}_{t+1}$ is a collection of ϵ-tight core constraints with respect to $\pi^{(t)}$. □

Lemma 4.3. *The Algorithm 1 terminates in at most $|N| = n$ iterations.*

Proof. By using Lemma 4.2, Algorithm 1 terminates, as $\mathcal{G}_t \neq \mathcal{G}_{t+1}$ for $t \geq 1$ and there are only finitely many such collections. Moreover, following the fact that $S_t \setminus T_t \neq \emptyset$, one can conclude $S_{t+1} \subsetneq S_t$. $S_0 = N$, and hence the algorithm terminates in at most $|N|$ iterations. □

Theorem 4.4. *Algorithm 1 returns an ϵ-bidder optimal core point, and uses $O(n \log(n/\epsilon))$ calls to the oracle* `WIN-ORAC`.

Proof. Combining Lemma 4.2 and Lemma 4.3, after at most $|N| = n$ iterations the algorithm terminates. At termination time T, $S_{T+1} = \emptyset$ is the subset of bidders that are participating simultaneously in all of the core constraints in the collection $\mathcal{G}_{T+1}$. Also, $\mathcal{G}_{T+1}$ is a collection of ϵ-tight core constraints with respect to $\pi^{(T)}$. Combining those, we conclude that the intersection of all ϵ-tight core constraints with respect to $\pi^{(T)}$ is empty. Now, $\pi^{(T)}$ is in the core and if you increase one of its coordinates, lets say j, by more than ϵ then we know there exists at least one ϵ-tight core constraint S with respect to $\pi^{(T)}$ such that $j \notin S$ and therefore by this change this constraint will be violated. So, by Definition 2.4, $\pi^{(T)}$ is ϵ-bidder optimal. Moreover, at each iteration t the algorithm uses one query call to `WIN-ORAC` to find S_{t+1}, and at most $\log(\frac{|S_{t+1}|}{\epsilon})$ query calls to do the binary search, so in total it only needs

$$\text{Total \# of oracle calls} \leq n + \sum_{t=0}^{T} \log\left(\frac{|S_{t+1}|}{\epsilon}\right) \leq n + \sum_{t=1}^{n} \log\left(\frac{t}{\epsilon}\right) = n + \log\left(\frac{n!}{\epsilon^n}\right) = O(n \log(n/\epsilon)).$$

5 Discussion

In this section, we elaborate on the computational and fairness properties of our core-selecting algorithm in comparison to existing payment rules in the literature, and then discuss the application of our results to online advertising.

Water-filling vs. VCG. VCG is a combinatorial auction in which reporting true values is a dominant strategy equilibrium. Computing payments in this auction can be done with $O(n)$ query calls to a optimal welfare oracle.[4] In comparison, our water-filling algorithm (Algorithm 1) induces a truncated strategy (i.e. $i \in N : b_i(.) = \max(v_i(.) - \pi_i, 0)$) that is a full-information ϵ-Nash equilibrium, as its payment rule is an ϵ-bidder optimal core payment [12]. Moreover, as we showed in Theorem 4.4, computing payments only requires $\tilde{O}(n)$ query calls to truncated maximum welfare allocation oracle. Furthermore, our payments produce more revenue at equilibrium [5] [10] and no coalition (subset of all bidders) can form a mutually beneficial renegotiation among themselves [10].

Parametrizing the Path and Fairness. The choice of direction for water-filling at each iteration t of Algorithm 1 is flexible, as long as it is only restricted to increasing utilities for bidders in set S_t. Therefore it can potentially implement different core payment rules tailored for different economical objectives. E.g., uniform water-filling will guarantee the approximate-equity of utilities subject to being a bidder optimal core point. Another desired objective is fairness with respect to VCG (e.g. [12] minimizes ℓ_1 distance to VCG, and [9] minimizes ℓ_2 distance to VCG). For this objective, we can consider a variant of our algorithm that at each iteration performs `Core-Search` along the ray that connects $\pi^{(t)}$ to VCG (See Figure 2). This VCG-pursuit heuristic finds a bidder optimal point that attempts to minimize the angle with the ray connecting the origin with the VCG outcome, and therefore it heuristically implements an equilibrium in which winning bidders receive (almost) the same fraction of their utilities in VCG.

Applications for Sale of Space Online Advertising Auctions. Online ads with multiple configurations (i.e., decorations) are becoming the most prominent way of displaying advertisements on the Web. Designing ad auctions is more complicated in this combinatorial setting, and standard solutions are either ill-defined (such as the Generalized Second Price (GSP) auction) or generate low revenue (such as VCG). One solution is to use core selecting auctions that are implementable at equilibrium [16]. However, the main drawback of core-selecting algorithms in the literature is either they are moderately fast heuristics (still much slower than quasi-linear oracle complexity) with no theoretical guarantees for convergence [12, 9, 14, 6], or are complex and slow convex programming algorithms. On the contrary, our proposed payment rule is significantly faster and uses at most $\tilde{O}(n)$ oracle calls. Note also that in such a setting one may be able to find the welfare maximizing allocation efficiently, or one can replace the optimal allocation with an approximate Maximal-In-Range allocation (see [13]) and get the same result as in this paper, but just for the core polytope restricted to that range.

[4] By relaxing the incentive constraints to hold in expectation, approximate VCG payments can be computed implicitly by just one call to the allocation oracle in single dimensional [4] or combinatorial [25].

[5] To compare the revenue, one can consider the coalition $N \setminus \{i\}$ and note that $\pi_i \leq w(N) - w(N \setminus \{i\})$, which is indeed the VCG utility of player i. Therefore, total generated revenue is lower-bounded by the VCG revenue.

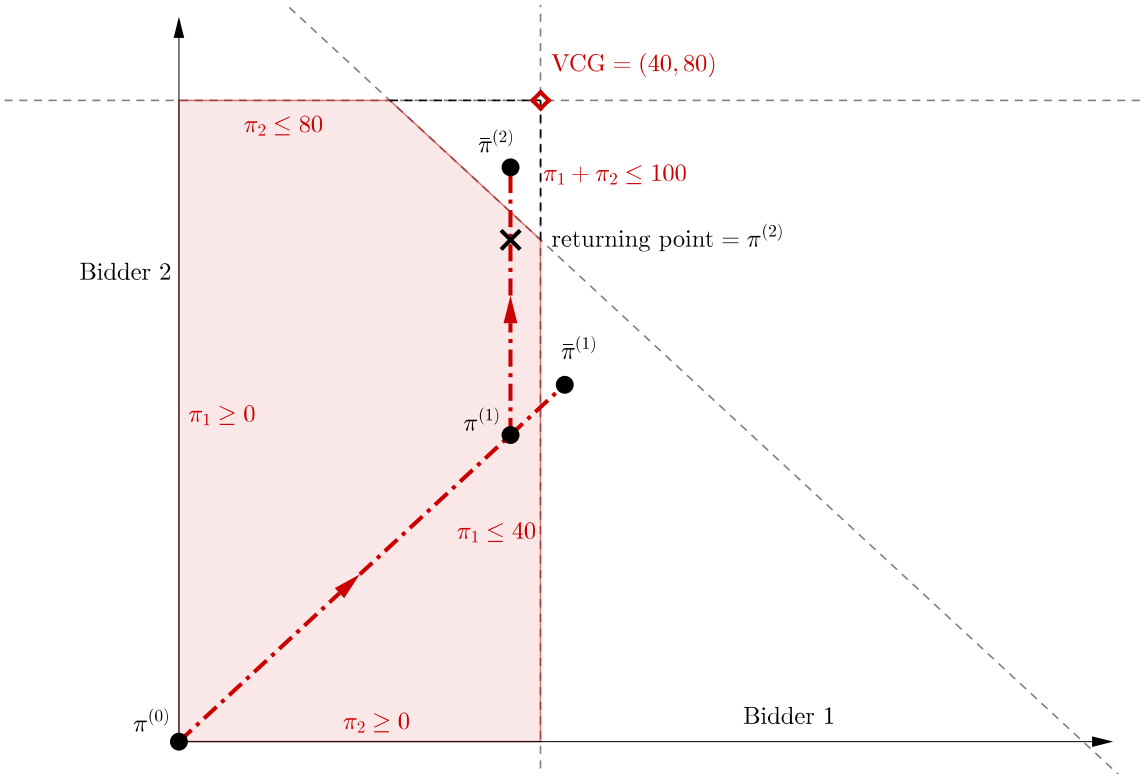

Figure 1: Water-filling algorithm for the following example in [12]. Let $m = 2$ items, A and B, $n = 5$ bidders, and let bids be as follow (each bidder submits one bid): $b_1(A) = 60$, $b_2(B) = 100$, $b_3(AB) = 60$, $b_4(A) = 20$, and $b_5(B) = 20$. The filled area is the core-polytope, and the arrow shows the path that algorithm follows. Note that $S_0 = \{1, 2, 3, 4, 5\}$, $T_0 = \{1, 2\}$, $S_1 = \{1, 2\}$, $T_1 = \{2, 4\}$, $S_2 = \{2\}$, $T_2 = \{3\}$ and $S_3 = \emptyset$.

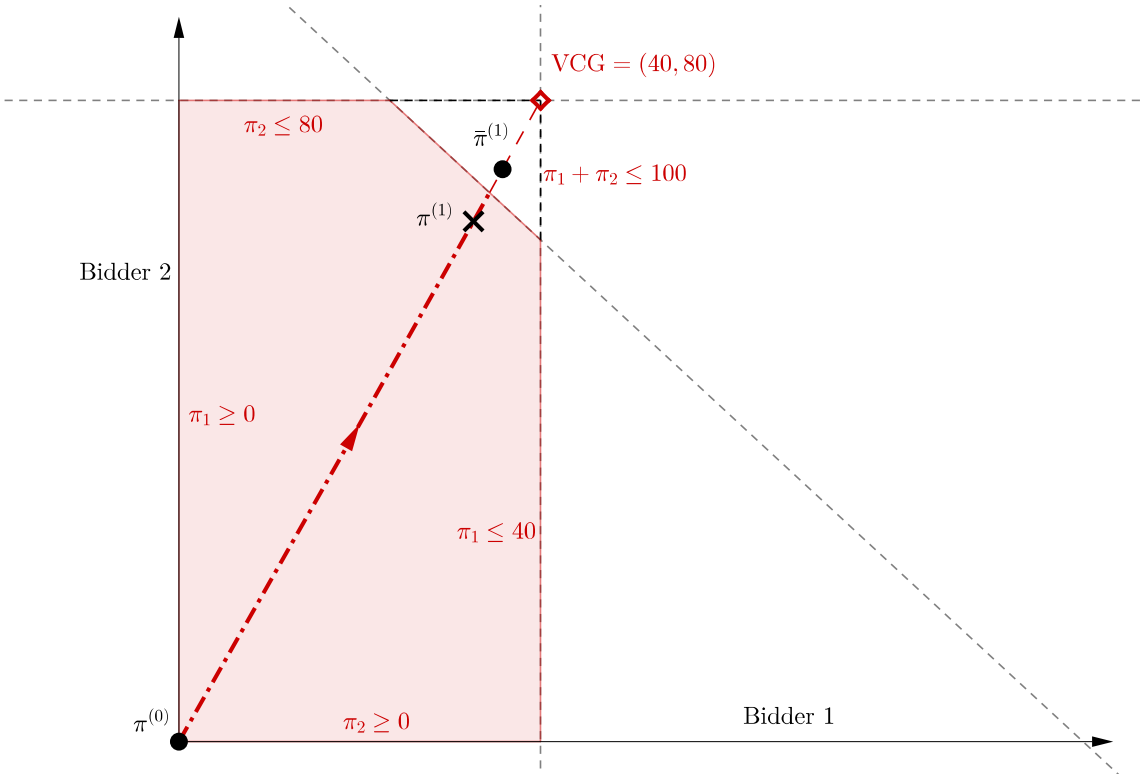

Figure 2: VCG-Pursuit

References

[1] L. M. Ausubel and O. V. Baranov. Core-selecting auctions with incomplete information. *University of Maryland*, 121, 2010.

[2] L. M. Ausubel and P. Milgrom. Ascending auctions with package bidding. *Advances in Theoretical Economics*, 1(1), 2002.

[3] L. M. Ausubel and P. Milgrom. The lovely but lonely vickrey auction. *Combinatorial auctions*, 17:22–26, 2006.

[4] M. Babaioff, R. D. Kleinberg, and A. Slivkins. Truthful mechanisms with implicit payment computation. In *Proceedings of the 11th ACM conference on Electronic commerce*, pages 43–52. ACM, 2010.

[5] M. Beck and M. Ott. Revenue monotonicity in core-selecting package auctions. Technical report, Working paper, 2009.

[6] B. Bünz, S. Seuken, and B. Lubin. A faster core constraint generation algorithm for combinatorial auctions. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 827–834. AAAI Press, 2015.

[7] E. H. Clarke. Multipart pricing of public goods. *Public choice*, 11(1):17–33, 1971.

[8] P. Cramton. Spectrum auction design. *Review of Industrial Organization*, 42(2):161–190, 2013.

[9] R. W. Day and P. Cramton. Quadratic core-selecting payment rules for combinatorial auctions. *Operations Research*, 60(3):588–603, 2012.

[10] R. W. Day and P. Milgrom. Core-selecting package auctions. *international Journal of game Theory*, 36(3-4):393–407, 2008.

[11] R. W. Day and P. Milgrom. Optimal incentives in core-selecting auctions. *Handbook of Market Design*, 2010.

[12] R. W. Day and S. Raghavan. Fair payments for efficient allocations in public sector combinatorial auctions. *Management Science*, 53(9):1389–1406, 2007.

[13] S. Dobzinski and N. Nisan. Mechanisms for multi-unit auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 346–351. ACM, 2007.

[14] A. Erdil and P. Klemperer. A new payment rule for core-selecting package auctions. *Journal of the European Economic Association*, 8(2-3):537–547, 2010.

[15] G. Goel and M. R. Khani. Revenue monotone mechanisms for online advertising. In *Proceedings of the 23rd international conference on World wide web*, pages 723–734. ACM, 2014.

[16] G. Goel, M. R. Khani, and R. P. Leme. Core-competitive auctions. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 149–166. ACM, 2015.

[17] J. K. Goeree and Y. Lien. On the impossibility of core-selecting auctions. *Theoretical economics*, 11(1):41–52, 2016.

[18] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981.

[19] T. Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631, 1973.

[20] L. Lamy. Core-selecting package auctions: a comment on revenue-monotonicity. *International Journal of Game Theory*, 39(3):503–510, 2010.

[21] D. C. Parkes, J. R. Kalagnanam, and M. Eso. Achieving budget-balance with vickrey-based payment schemes in exchanges. 2001.

[22] B. Rastegari, A. Condon, and K. Leyton-Brown. Revenue monotonicity in deterministic, dominant-strategy combinatorial auctions. *Artificial Intelligence*, 175(2):441–456, 2011.

[23] R. Sano. Non-bidding equilibrium in an ascending core-selecting auction. *Games and Economic Behavior*, 74(2):637–650, 2012.

[24] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

[25] C. A. Wilkens and B. Sivan. Single-call mechanisms. *ACM Transactions on Economics and Computation*, 3(2):10, 2015.